

Future Technology Challenges for Collaboration: Customizing the Grid

Position Paper for the Department of Energy

Jon B. Weissman

Department of Computer Science and Engineering

University of Minnesota, Twin Cities

Minneapolis, MN 55455

Voice: (612) 626-0044

Fax: (612) 625-0572

WWW: www.cs.umn.edu/~jon

E-mail: jon@cs.umn.edu

Introduction

Collaboration is the spatial and temporal integration of disparate skills and resources for the purpose of solving a problem that would otherwise be difficult or impossible to solve. However, collaborations don't always succeed or get initiated for a multitude of reasons, both technical and sociological, including: insufficient collective incentive to working together, high "latency" to get started due to a lack of "clean" problem decomposition and/or interfacing, and "roll-your-own" cultures that are resistant to change. This is unfortunate in that successful collaborations are often "greater than the sum of their parts". In this position paper, we address the "roll-your-own" issue specifically focusing on the integration and use of software provided by external collaborators. The hesitancy to use outside software may slow down a collaboration or prevent its formation. For example, consider two groups, group A supports computational and analysis tools, and group B supports visualization tools. Further suppose that the scientists in group A and B would like to work together on common problems involving computation, analysis and visualization. One conduit for this collaboration could be the interaction enabled by the shared tools provided by each group.

We speculate that the "roll-your-own" attitude towards the use of software provided by collaborators or potential collaborators stems in part from inertia but also from the belief that such software will not exactly match the needs of the group. They are not incorrect. High-performance Grid software, like most software, is designed for a single purpose with a single objective (defined by its implementors). In the new and emerging world of Grid services, there is an opportunity: by designing services with flexibility in mind, from the outset, the barriers to collaboration enabled by the use of shared services may be lowered.

Flexible Services

A *flexible* Grid service exposes "dials" that can be turned to meet the specific Quality-of-Service (QoS) needs of its users through customization. Dials may relate to (i) *performance* features, e.g. service throughput or min/max completion time, (ii) *reliability* features, e.g. number of spare resources for service request execution, whether checkpointing is on or off, the number of "9's" for success, (iii) *security* features, e.g. encrypt all traffic to and from the service, or (iv) *functionality* features, e.g. use one of Blast, Smith-Waterman, or Fasta algorithm for the gene

sequence comparison algorithm. Service customization raises several challenging technical issues:

- *how can service flexibility be implemented and what is the incentive for the service provider to provide it?*
- *how does the user or client notion of QoS map to the various dial settings?*
- *how can competing customizations to the same service be handled (e.g. user A and user B want conflicting features from the same service)? is there a useful notion of "best effort" QoS? can we guarantee some kind of "QoS" isolation?*
- *how can feature dependence be represented and handled (e.g. turning the dial on a reliability dial may come at the cost of performance)?*
- *how can service discovery be enhanced to include customized services (e.g. a client may wish to locate a genomic service which provides a performance dial that can deliver a throughput between X and Y)?*
- *can composite services (a service that contains or uses other services) be made flexible? this applies to workflow pipelines.*
- *to what extent does the dynamics and unpredictable nature of the Grid impact customization. are QoS guarantees merely stochastic?*
- *what common middleware could be provided in support of flexible services?*

The development of flexible services will require a degree of engineering that goes beyond current Grid service deployments. However, we feel that future production-scale Grid collaborations of importance to DOE will greatly benefit from this paradigm.